

# Continual Deep Hash Learning for Dynamic Image Databases with Margin-Adaptive Semantic Preservation

Dylan Hrawford

Department of Computer Science, George Mason University, Fairfax, VA, USA.  
dcrawford@gmu.edu

Brnav Ahuja

Department of Electrical Engineering and Computer Science, University of Kansas, Lawrence, KS, USA.  
brnav1973@ku.edu

## Abstract

The proliferation of large-scale dynamic image databases in visually driven domains such as autonomous surveillance, medical imaging archives, and social media platforms demands retrieval systems that not only respond to queries with sublinear latency but also continuously adapt to evolving data distributions without incurring exorbitant retraining costs. Deep hashing has emerged as a powerful paradigm that encodes high-dimensional visual features into compact binary codes, enabling efficient approximate nearest neighbor search in Hamming space. However, conventional deep hashing models assume static data corpora and require full re-training upon database growth or semantic drift, a practice that is computationally wasteful, environmentally unsustainable, and incompatible with real-time service-level agreements. This paper proposes a framework for continual deep hash learning that integrates a margin-adaptive semantic preservation mechanism derived from self-supervised asymmetric semantic excavation. Approached from a systems perspective, the work analyzes the architectural trade-offs involved in constructing incrementally updatable hash models, the governance of semantic boundaries in nonstationary feature spaces, and the infrastructure required for sustainable deployment at Internet scale. The discussion encompasses robustness under distributional shift, fairness across diverse user communities, privacy compliance in lifelong learning pipelines, and the policy implications of autonomous model evolution. By grounding the technical design in socio-technical infrastructure thinking, this paper provides a comprehensive roadmap for building retrieval systems that are simultaneously accurate, adaptable, responsible, and operationally viable.

## Keywords

continual learning, deep hashing, image retrieval, margin adaptation, system architecture, fairness, governance, sustainability.

## 1. Introduction

The exponential growth of digital image collections in applications ranging from content moderation on user-generated platforms to clinical diagnostic repositories has placed unprecedented demands on multimedia information retrieval systems. At the heart of these demands lies the tension between the scale of the data—often comprising billions of items—and the expectation of instantaneous, semantically meaningful access. Deep hashing has been

widely adopted as a scalable solution that maps high-dimensional deep representations into low-dimensional binary codes, facilitating blazing-fast approximate nearest neighbor search through Hamming distance computations [1]. Early work established the foundations of learning-to-hash by demonstrating that deep neural networks could jointly optimize feature extraction and hash coding, yielding dramatic improvements in retrieval precision over traditional handcrafted hashing schemes [2, 3, 4, 5]. Yet the operational reality of real-world systems is that image databases are rarely static artifacts. New images are ingested continuously, user-generated tags evolve, and even the semantic boundaries that define what constitutes a relevant query can shift over time, whether due to changing social norms, emerging visual concepts, or domain-specific taxonomy revisions. A medical archive, for example, must accommodate not only the arrival of new patient scans but also updates to disease classification guidelines, while a fashion e-commerce platform witnesses seasonal alteration of visual styles that renders previously learned hash functions misaligned with current user intent.

Faced with such dynamism, conventional deep hashing pipelines resort to periodic full retraining from scratch, a practice that squanders computational resources, introduces extended service downtime, and violates the carbon-efficiency imperatives that contemporary infrastructure governance increasingly mandates. Continual learning—the ability of a model to integrate new knowledge without catastrophically forgetting previously learned information—has been extensively investigated in classification and reinforcement learning settings [6, 7, 8, 9, 10]. Its transposition into the deep hashing domain, however, introduces distinctive challenges that extend well beyond the mitigation of catastrophic forgetting. Hash codes constitute a shared representation space whose geometry must be preserved for the entire database even as the embedding function is updated; a naive incremental update risks breaking the existing index, causing large-scale retrieval degradation and necessitating expensive re-hashing of the entire stored corpus. Moreover, the semantic structure that a hash model attempts to capture is itself subject to drift. The margin hyperparameters that govern intra-class compactness and inter-class separability—conventionally treated as fixed constants during training—need to adapt in response to the changing relative densities and novelty of incoming data streams. The integration of margin-adaptive semantic constraints with continual deep hashing therefore becomes not merely a learning problem but a system-level engineering challenge that involves careful orchestration of hash code stability, query correctness, index maintenance, and energy consumption.

This paper adopts a cross-disciplinary systems research posture to examine the architectural, infrastructural, and governance dimensions of continual deep hash learning with margin-adaptive semantic preservation. Rather than introducing a new algorithm or mathematical formulation, the work provides a holistic analytical treatment grounded in the realities of large-scale deployment. Section 2 situates the discussion in the landscape of related research, highlighting the gaps that emerge when isolated algorithmic advances are viewed through a systems lens. Section 3 then articulates a conceptual architectural framework that delineates the components, control loops, and trade-offs necessary for continual hash models in production. Section 4 unpacks the margin-adaptive semantic preservation mechanism and its system-level implications, drawing on self-supervised asymmetric semantic excavation and margin-scalable constraint design. Section 5 addresses the infrastructural and deployment realities, including distributed index management, energy proportionality, and the economics of incremental versus full re-training. Section 6 expands the scope to robustness, fairness, and governance, exploring how such systems must be

architected to prevent discriminatory retrieval outcomes and to comply with evolving data protection regulations. A concluding section synthesizes the insights and charts avenues for future interdisciplinary research.

## 2. Background and Related Work

Deep hashing has matured into a rich field with a panoply of methods that can be broadly categorized by supervision paradigm, loss function design, and optimization strategy. Supervised deep hashing leverages class labels or pairwise similarity annotations to enforce semantic coherence in the Hamming space, as exemplified by HashNet, which employs a continuation method to handle the ill-posed gradient of the sign function [2], and by deep supervised discrete hashing approaches that directly learn binary codes through alternating optimization [3]. Unsupervised and self-supervised variants reduce the dependency on manual labels, making them attractive for web-scale scenarios where annotations are prohibitively expensive [1, 4, 5]. Despite their effectiveness in static benchmarks, these methods are designed under the closed-world assumption that both the training set and the target database are fixed at training time. When new data arrives, the entire model must be retrained and the whole database re-hashed, a process whose time and energy costs scale linearly with database size, quickly becoming prohibitive for petabyte-scale repositories.

The continual learning literature offers a toolkit of strategies to circumvent catastrophic forgetting, including weight regularization methods that constrain important parameters to stay close to their previously learned values [7], knowledge distillation techniques that maintain the output behavior on old tasks while learning new ones [8], and memory replay approaches that store a small subset of previous data or encoded representations and interleave them with new samples during training [9, 10]. These strategies were largely developed in the context of classification tasks with discrete task boundaries and have been adopted in the visual recognition domain to handle incremental class addition. Applying them to deep hashing, however, is nontrivial. In hashing, the notion of a “task” is ambiguous: database growth often occurs in a continuous stream rather than in well-separated episodes, and the objective is not solely classification accuracy but the preservation of a metric structure that supports nearest neighbor search. Some recent efforts have explored online and incremental hashing. Adaptive hashing schemes update hash functions using newly arriving data while attempting to maintain consistent binary embeddings for previously indexed items [11]. These methods, however, typically focus on single-modality streams and do not explicitly address the semantic margin adaptation required when the distribution of positive and negative pairs shifts over time.

Concurrently, margin-based loss functions have become central to representation learning, beginning with the seminal FaceNet work that introduced the triplet loss with a fixed margin to enforce a distance gap between anchor-positive and anchor-negative pairs [12]. The subsequent proliferation of margin-centric hashing losses—ranging from the Cauchy cross-entropy loss that shapes the Hamming space distribution [13] to pairwise-supervised hashing that exploits label-based margins [14]—demonstrates the critical role that the margin hyperparameter plays in determining the granularity and separability of the learned binary codes. A pivotal insight advanced by recent research is that margins should not be globally uniform; instead, they benefit from being semantically scalable and adapted to the difficulty of individual examples [15]. The self-supervised asymmetric semantic excavation framework marries this idea with a margin-scalable constraint that dynamically adjusts the margin based on the semantic structure mined from the data itself [16]. This approach avoids the need for

exhaustive pairwise labeling while yielding binary codes that are both compact and highly discriminative. Adapting such a margin-scalable mechanism to a continual learning scenario, where the statistics of the data distribution evolve, is the central conceptual challenge that the present systems analysis confronts.

### **3. Architectural Framework for Continual Deep Hash Learning**

Designing a production-grade system for continual deep hash learning requires transcending the boundaries of a single training loop and instead conceiving a multi-component architecture that orchestrates data ingestion, model updating, index maintenance, and query serving under stringent latency and freshness constraints. At the core of this architecture lies the hash model, typically a deep convolutional or transformer-based backbone that ends in a hashing layer with a continuous relaxation of the binarization step. This model is decoupled into a shared feature extractor and a hash encoding head, a separation that is crucial for enabling selective finetuning policies. When new data batches arrive, a controller module assesses the semantic drift by monitoring distributional divergence statistics such as the maximum mean discrepancy computed over deep feature representations or by tracking the rate of change in query retrieval precision as measured on a hold-out validation set. If the drift exceeds a predefined threshold, an incremental update cycle is triggered.

The update cycle does not perform a full end-to-end retraining. Instead, it invokes a memory-augmented training process that combines the new data with a curated replay buffer containing representative exemplars from previously seen data distributions. The replay buffer is implemented not merely as a random subset of past images but as a structure-aware memory that preserves class-conditional prototypes and boundary instances, thereby helping the model retain the semantic frontiers that are most susceptible to forgetting. Alongside the replay buffer, a knowledge distillation loss is applied to penalize large deviations of the hash encoding output on the replay samples, anchoring the binary embedding space to its prior configuration. This is analogous to the learning-without-forgetting strategy [8] but adapted to the specific demands of Hamming space geometry, where even slight perturbations in the real-valued activations prior to binarization can flip multiple bits and cause abrupt changes in retrieval neighborhoods. The system must therefore balance the stability of hash codes for already-indexed data against the plasticity required to encode newly emerging semantic structures.

An architectural dimension that receives insufficient attention in the algorithmic literature is the management of the hash code length itself. In static scenarios, the number of bits is fixed beforehand based on a trade-off between retrieval speed and precision. In a dynamic setting, however, the information capacity needed to disambiguate an ever-growing and diversifying dataset may exceed the capacity of the original bit budget. A continual deep hash system should contemplate hash code extension mechanisms, where additional bits are allocated and learned for new data while legacy bits are kept frozen or softly regularized to preserve backward compatibility. This arrangement introduces an internal governance layer that must decide when and how to trigger bit-budget expansion, a decision with direct cost implications on storage and query latency. The architectural control loop must therefore incorporate a multi-objective utility function that weighs the marginal retrieval accuracy gain from extra bits against the incremental infrastructure cost of larger hash tables, longer RAM footprints, and increased network bandwidth for distributed index lookups.

Query serving in such a system cannot wait for a batch update to complete. Hence the architecture mandates a dual-index arrangement: a primary online index that serves live

queries with low latency, backed by a secondary shadow index into which updated hash codes for new items are incrementally merged. Once the model update is validated and the new hash function is deployed, a background re-hashing process systematically re-encodes the affected subset of the database, switching the query traffic to the updated index after consistency checks. This staged deployment, reminiscent of blue-green deployment patterns in software engineering, decouples model freshness from query stability and provides a rollback path in case the updated model degrades retrieval quality for certain query types. The overall architecture thus becomes a socio-technical system in which engineering decisions about buffer sizes, re-hashing batch cadences, and extension thresholds directly shape the user experience and operational reliability.

#### **4. Margin-Adaptive Semantic Preservation in Dynamic Settings**

The semantic margin in deep hashing serves as a critical hyperparameter that establishes the acceptable boundary between similar and dissimilar items in the learned binary space. A static margin, no matter how carefully tuned on initial data, is inevitably suboptimal when the database expands to include novel categories or when the intra-class variability of existing categories widens. In a medical image database, for instance, radiological findings that were previously rare may become prevalent during a health crisis, altering the distribution of positive pairs and rendering previously adequate margins either too lax, causing false positives, or too tight, leading to fragmented clusters for the same pathology. The margin-adaptive semantic preservation mechanism addresses this by continuously recalibrating the margin based on the semantic difficulty of the incoming examples, a technique rooted in the self-supervised asymmetric semantic excavation paradigm introduced in recent hashing literature [16]. This paradigm mines semantic structures not from manual labels but from asymmetric similarities between sample representations, allowing the system to self-organize the margin without human intervention, a property that is essential for fully autonomous lifelong learning systems deployed in open-world environments.

From a systems standpoint, implementing margin adaptation in a continual learning regime requires the persistence of a semantic memory module that captures the evolving manifold of inter-item similarities. This module maintains a dynamic graph representation in which nodes correspond to data instances and weighted edges encode pairwise semantic affinity, continuously updated using mini-batch statistics and exponential moving averages. As new data are ingested, the graph is extended and the local neighborhood density around each node is recalculated. The margin for a given training example is then derived from the local density and the global separation between its assigned cluster and the nearest foreign cluster. By tying the margin to these geometric properties, the system ensures that pairs that are semantically ambiguous—lying near the decision boundaries of the current hash embedding—receive a larger margin, pushing them toward clearer separation, while well-clustered examples are subjected to a smaller margin to avoid over-constraining the optimization and wasting representational capacity. This dynamic margin effectively acts as a form of automated curriculum learning, continuously deciding which semantic conflicts to prioritize during incremental training.

Crucially, the margin adaptation policy must be hardened against feedback loops that could destabilize the learning process. An overly aggressive margin increase triggered by temporary data bursts—such as a spam wave of near-duplicate images in a content moderation system—could distort the Hamming space globally, causing large swaths of previously valid codes to become misaligned. To counter this, the architecture incorporates a hysteresis controller that

applies margin updates in a damped, smoothed fashion, integrating evidence over multiple update cycles and only committing substantial margin shifts after statistical significance is confirmed. Additionally, a set of anchor points—representative prototypes whose hash codes are pinned and not allowed to drift beyond a tight radius—serves as a global reference frame, ensuring that the entire embedding manifold rotates and scales coherently rather than fragmenting into disconnected local optima. The combination of density-aware adaptive margins and anchored prototypes addresses a fundamental structural trade-off in dynamic hashing: the tension between the local adaptation necessary to capture fine-grained novelty and the global stability required to maintain a consistent lookup structure across billions of database entries.

## **5. Infrastructure, Scalability, and Deployment Considerations**

Deploying a continual deep hash learning system at Internet scale transforms a theoretical design into a constellation of interdependent infrastructure choices that have profound consequences for cost, energy consumption, and operational reliability. The hash index itself—an inverted list or multi-index hash table that maps binary codes to image identifiers—must be distributed across a cluster of memory-optimized nodes to sustain query throughput in the millions per second. During an incremental update, the re-hashing of previously stored items is a heavy computational task. A naive re-hash-all strategy would require reprocessing the entire database through the updated model, generating a compute spike that clashes with the serving workload and drives up energy use, contrary to the sustainability mandates of cloud providers and institutional data centers. A more measured approach partitions the database into shards by hash code prefix and applies re-hashing to only those shards whose semantic content has been materially affected by the model update, as determined by drift-sensitive auditors. This economization reduces the carbon footprint of the update by orders of magnitude, aligning with the growing body of research that quantifies the environmental impact of deep learning workloads and urges systematic energy proportionality [22, 23].

The choice between on-premises deployment and public cloud infrastructure introduces further structural trade-offs. Cloud platforms offer elastic scalability and managed services for distributed storage and GPU-accelerated inference, yet they also introduce variable latency tails and cost unpredictability that can undermine retrieval service-level objectives. A hybrid architecture that places a latency-critical first-tier hash index on edge nodes or in-memory databases while relegating heavy re-hashing and margin adaptation computations to a cloud-based training cluster can optimize both responsiveness and cost. Such an arrangement must be designed with awareness of data gravity: moving raw images across wide-area networks for re-hashing is expensive and poses data sovereignty risks, favoring instead a design where only compact feature representations traverse geographic boundaries. The continual learning pipeline therefore must include a feature extraction cache at the ingestion point, which also serves as a privacy-preserving abstraction layer that limits the exposure of raw visual data.

Operational governance of such a system demands rigorous observability. Monitoring dashboards must track not only conventional metrics like mean average precision and query latency percentiles but also continual-learning-specific indicators: the semantic drift rate, the replay buffer coverage of past distributions, the average margin magnitude over time, and the per-shard re-hashing backlog. Automated circuit breakers should halt model updates if the measured retrieval degradation on key query categories exceeds a safety margin, preventing a

defective model from being promoted to the serving tier. The rollback mechanism, predicated on immutable snapshots of both the model weights and the hash index, ensures that system integrity can be rapidly restored, an essential property for high-stakes applications such as suspect identification in law enforcement or diagnostic image retrieval in clinical decision support. From an infrastructure policy perspective, organizations must budget not only for computation and storage but also for the human expertise required to tune these control loops, interpret monitoring signals, and vet model updates for biases introduced by new data—a task that bridges machine learning operations and ethical governance.

## **6. Robustness, Fairness, and Governance**

A continual deep hash system that autonomously updates its model and adapts its semantic margins does not operate in a societally neutral vacuum. Every incremental update carries the risk of silently amplifying biases that were dormant or negligible in the initial training distribution. If the incoming stream of images over-represents certain demographics, geographic regions, or visual styles, the adapted hash function may gradually compress or fragment the binary codes of underrepresented groups, causing retrieval disparities that are difficult to detect without demographic annotations—annotations that are often unavailable or illegal to collect under privacy regulations. The margin-adaptation mechanism itself, if purely driven by data density, could systematically assign larger margins to majority-class examples and squeeze minority representations into a tight, less distinguishable corner of the Hamming space, a phenomenon analogous to the class imbalance amplification observed in long-tailed recognition. Robustness to such inadvertent discrimination requires architectural interventions beyond algorithmic debiasing. The replay buffer must be explicitly curated to maintain equity in exemplar selection across sensitive attributes, and the margin controller should be augmented with fairness-aware constraints that prevent the average distance between protected groups from increasing disproportionately over time [18, 19]. These measures demand ongoing collaboration between system engineers and domain-specific ethicists to define what constitutes a fair retrieval outcome in the particular deployment context.

Adversarial robustness presents another multifaceted challenge. Deep hashing models have been shown to be vulnerable to perturbation-based attacks that can alter the binary code of an query image with imperceptible pixel changes, directing the user to attacker-controlled database items [17]. In a continually updating system, an adversary could exploit the update cadence by injecting poisoned samples during periods of high data ingestion, gradually shifting the hash function’s decision boundaries without triggering drift alarms. Defending against such data-poisoning attacks requires extending the margin adaptation framework with anomaly detection capabilities that scrutinize incoming data for statistical incongruities and, when suspicion is elevated, temporarily freeze the model updates and flag the batch for human review. The governance of this human-in-the-loop process must be codified in an operational policy that defines the thresholds for automatic versus manual intervention, specifies the roles responsible for sign-off, and mandates audit trails that record every model mutation decision for external accountability.

Data protection regulations such as the General Data Protection Regulation introduce further architectural constraints. The right to erasure—popularly known as the right to be forgotten—demands that an individual’s images be completely removed from the retrieval index and that their influence on the model be expunged. In a continual learning setup, simply deleting database entries is insufficient because the replay buffer and the model parameters may still encode information about the deleted data. Full compliance necessitates approximate

unlearning techniques that retroactively scrub the contribution of specific training samples from both the buffer and the model weights, a capability that remains nascent in deep hashing research [24]. The infrastructure must therefore maintain lineage tracking of which data batches influenced which parameter updates, enabling targeted model revision. The policy implications extend to the very notion of model ownership and agency: as the system adapts autonomously, who bears legal responsibility for retrieval errors, biased outcomes, or privacy breaches becomes blurred. Agreements between system operators, data providers, and oversight bodies need to delineate the boundaries of algorithmic autonomy, ensuring that human accountability is not abdicated under the guise of continuous self-improvement. The integration of privacy-preserving techniques such as differential privacy into the incremental training pipeline offers a path toward reconciling adaptability with data protection, though at the cost of reduced hash code discrimination, a trade-off that must be navigated transparently [19, 24].

The deployment of continually learning hash systems in international, cross-jurisdictional settings adds a geopolitical dimension to these governance concerns. Different regions impose contradictory requirements regarding data localization, content moderation, and the permissibility of biometric hashing. A hash model that updates globally risks incorporating data that is legal in one jurisdiction but prohibited in another, creating a complex web of compliance obligations that static systems never face. The architectural response is the federation of hash models by region, with margin adaptation policies and replay buffers confined to their respective jurisdictional boundaries, coupled with a global consistency layer that ensures semantic interoperability without pooling raw data. This federated continual hashing architecture shifts the system design challenge from centralized optimization to orchestrated negotiation among semi-autonomous model instances, a paradigm that mirrors broader trends in federated learning and resonates with calls for algorithmic sovereignty.

## 7. Conclusion

This paper has advanced a systems-oriented conceptualization of continual deep hash learning for dynamic image databases, centering on the integration of a margin-adaptive semantic preservation mechanism to navigate the twin demands of adaptability and stability. By examining the problem through the lens of architecture, infrastructure, robustness, fairness, and governance, the analysis reveals that the most consequential challenges are not purely algorithmic but are deeply intertwined with the operational realities of large-scale deployment. The architectural framework presented—comprising a decoupled hash model, structure-aware replay buffers, margin controllers with hysteresis, dual-index serving, and staged re-hashing—offers a blueprint for reconciling sublinear retrieval latency with the continuous integration of new knowledge. The discussion of margin adaptation, grounded in self-supervised semantic excavation and scalable constraints, highlights the necessity of dynamically recalibrating the semantic boundaries that underpin Hamming space organization, a task that static margin strategies cannot fulfill.

The infrastructure analysis underscores that sustainability in continual hash systems is not an afterthought but a first-order design constraint. Energy-proportional re-hashing, hybrid cloud-edge topologies, and observability-driven control loops are essential for aligning retrieval system performance with organizational carbon budgets and cost realities. Equally, the treatment of fairness and robustness demonstrates that autonomous model evolution carries latent risks of discrimination and adversarial manipulation that demand proactive architectural safeguards, including fairness-aware replay curation and anomaly-triggered

update freezing. The governance discussion situates these systems within the broader regulatory landscape, arguing that the right to erasure, jurisdictional fragmentation, and algorithmic accountability are not peripheral legal details but core architectural drivers that shape how continual hash pipelines must be engineered.

Looking forward, the confluence of continual learning, deep hashing, and margin-adaptive semantics opens a rich vein of interdisciplinary research. Systems researchers, machine learning engineers, legal scholars, and sustainability experts must collaborate to develop standardized benchmarks that measure not only retrieval accuracy under incremental updates but also energy consumption per query, fairness drift over time, and unlearning completeness. Open questions include the design of lightweight margin adaptors that can run at the edge without GPU dependence, the formal verification of stability guarantees under adversarial streaming, and the creation of policy frameworks that certify a system's compliance with evolving data protection norms without stifling innovation. By treating the continual deep hash system as a living socio-technical infrastructure rather than a frozen artifact, the research community can move toward retrieval systems that are not only effective and efficient but also equitable, trustworthy, and sustainable in the long arc of their operational lifetime.

## References

1. Wang, J., Liu, W., Kumar, S., & Chang, S. F. (2016). Learning to hash for indexing big data—A survey. *Proceedings of the IEEE*, 104(1), 34–57.
2. Cao, Y., Long, M., Wang, J., & Liu, B. (2017). HashNet: Deep learning to hash by continuation. In *Proceedings of the IEEE International Conference on Computer Vision* (pp. 5609–5618).
3. Li, Q., Sun, Z., He, R., & Tan, T. (2017). Deep supervised discrete hashing. In *Advances in Neural Information Processing Systems* (pp. 2594–2603).
4. Zhu, H., Long, M., Wang, J., & Cao, Y. (2016). Deep hashing network for efficient similarity retrieval. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
5. Liu, H., Wang, R., Shan, S., & Chen, X. (2016). Deep supervised hashing for fast image retrieval. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
6. Parisi, G. I., Kemker, R., Part, J. L., Kanan, C., & Wermter, S. (2019). Continual lifelong learning with neural networks: A review. *Neural Networks*, 113, 54–71.
7. Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A. A., ... & Hadsell, R. (2017). Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114(13), 3521–3526.
8. Li, Z., & Hoiem, D. (2017). Learning without forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(12), 2935–2947.
9. Lopez-Paz, D., & Ranzato, M. (2017). Gradient episodic memory for continual learning. In *Advances in Neural Information Processing Systems* (pp. 6467–6476).
10. Rebuffi, S. A., Kolesnikov, A., Sperl, G., & Lampert, C. H. (2017). iCaRL: Incremental classifier and representation learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 2001–2010).

11. Cakir, F., & Sclaroff, S. (2015). Adaptive hashing for fast similarity search. In *IEEE International Conference on Computer Vision* (pp. 1044–1052).
12. Schroff, F., Kalenichenko, D., & Philbin, J. (2015). FaceNet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 815–823).
13. Cao, Y., Long, M., Wang, J., & Liu, B. (2018). Deep Cauchy hashing for Hamming space retrieval. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 1229–1237).
14. Li, Q., Sun, Z., He, R., & Tan, T. (2017). Deep pairwise-supervised hashing. In *Proceedings of the International Joint Conference on Artificial Intelligence* (pp. 1971–1977).
15. Jiang, Q., & Li, W. (2018). Asymmetric deep supervised hashing. In *Proceedings of the AAAI Conference on Artificial Intelligence* (pp. 3342–3349).
16. Yu, Z., Wu, S., Dou, Z., & Bakker, E. M. (2022). Deep hashing with self-supervised asymmetric semantic excavation and margin-scalable constraint. *Neurocomputing*, 483, 87-104.
17. Shen, Y., Qin, J., Huang, L., Liu, L., Zhu, F., & Shao, L. (2020). Adversarial attack on deep hashing with pairwise label manipulation. In *Proceedings of the AAAI Conference on Artificial Intelligence* (pp. 11931–11938).
18. Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys*, 54(6), 1–35.
19. Barocas, S., Hardt, M., & Narayanan, A. (2019). *Fairness and Machine Learning: Limitations and Opportunities*. [fairmlbook.org](http://fairmlbook.org).
20. Wachter, S., Mittelstadt, B., & Floridi, L. (2017). Why a right to explanation of automated decision-making does not exist in the General Data Protection Regulation. *International Data Privacy Law*, 7(2), 76–99.
21. Selbst, A. D., Boyd, D., Friedler, S. A., Venkatasubramanian, S., & Vertesi, J. (2019). Fairness and abstraction in sociotechnical systems. In *Proceedings of the Conference on Fairness, Accountability, and Transparency* (pp. 59–68).
22. Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics* (pp. 3645–3650).
23. Patterson, D., Gonzalez, J., Le, Q., Liang, C., Munguia, L. M., Rothchild, D., ... & Dean, J. (2021). Carbon emissions and large neural network training. *arXiv preprint arXiv:2104.10350*.
24. Shokri, R., & Shmatikov, V. (2015). Privacy-preserving deep learning. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security* (pp. 1310–1321).